 By J Stoddard, MakeSans Systems, and A Goodrum, Compaq Computer Corp

AS PCI-X AND CONVENTIONAL, 66-MHz PCI DEVICES AND SYSTEMS INCREASINGLY APPEAR IN A RANGE OF INDUSTRY-STANDARD COMPUTER ARCHITECTURES, SYSTEM DESIGNERS CONFRONT NEW CHALLENGES IN DEVELOPING COST-EFFECTIVE PRODUCTS THAT COMPLY WITH THE PCI-ELECTRICAL SPECIFICATIONS.

Electrical-design considerations for PCI-X and 66-MHz PCI cards

In the past, most systems had sufficient design margin that they could violate PCI-electrical-specification parameters without causing failure. With PCI-X and conventional, 66-MHz PCI, however, data rates are increasing. Available device-signal edge rates are increasing even faster than data rates. Therefore, designers must fully comply with the specification to guarantee that their cards will work in all systems. This article examines the critical elements of the PCI-X and conventional 66-MHz PCI-electrical specifications, the problems you may encounter by using design shortcuts, and appropriate design methods for successful production of PCI-X and conventional PCI add-in cards. You need a solid understanding of high-frequency design techniques to properly manage electrical noise, such as crosstalk, ground bounce, and reflections.

THE I/O-BUFFER PROBLEM AND IC PACKAGES

The I/O buffer and the topology of the signal connections, including the host bridge and connectors, determine the shape of the signal waveforms. System designers are responsible for designing systems in which all the signals settle to their appropriate states within the specified input-setup time for the full range of device and add-in-card characteristics the specification allows. However, I/O buffers in many ASIC libraries drive harder and have faster edge rates than the PCI specifications allow. Many designers want to use these faster buffers to help meet the faster output-delay times that the PCI-X and conventional 66-MHz PCI specifications require. However, using devices that are stronger or faster than the specifications allow can cause excessive reflective noise, which causes glitches at device inputs after the signal should be stable. This noise causes the signal to violate the device-input setup-time requirement. Using these stronger or faster devices also causes increased crosstalk, which causes false transitions on quiet lines and erroneous triggering; increased ground and supply bounce, which causes false transitions on quiet lines and erroneous triggering; and higher noise levels on supply voltages, which increase electromagnetic-control problems and decrease signal margins. Reduced signal margins in turn may cause logic errors. To solve these problems, you should use I/O buffers with drive strengths and slew rates that meet PCI-X and PCI specifications.

Other concerns in designing for PCI-X and conventional PCI include the IC-package, supply voltage, and return paths. Improper design of these elements slows transition times, reducing setup times; may cause ringing, causing false double transitions and uncertain timing; and may cause coupling between signals, causing false transitions on quiet lines.

The underlying principle of good package performance is the intimacy between each signal and its

Figure 1

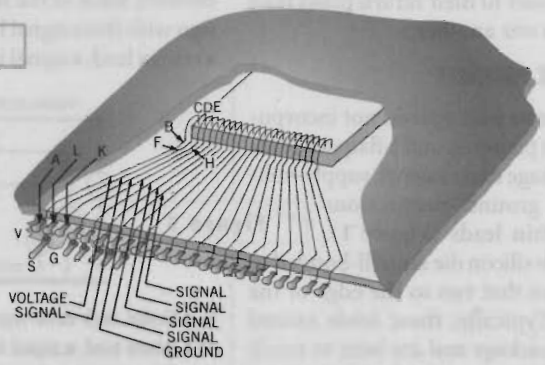

A quad-flatpack lead-frame package makes signal, supply-voltage, and ground connections through thin leads.

return path. Ideal packages or connectors position each signal and its return close together for two reasons. First, the closer a signal is to its return path, the heavier the coupling to that return and the lighter the coupling to other signals will be. Conversely, the farther a signal is from its return path, the larger the current loop is, the higher self-inductance and mutual inductance become, and the greater the likelihood that the signal will couple with other current loops. Second, the tighter the coupling between signal and return paths, the higher the maximum achievable operating frequency becomes. Designers should keep signals and returns close together as they pass through connectors, packages, or other discontinuities.

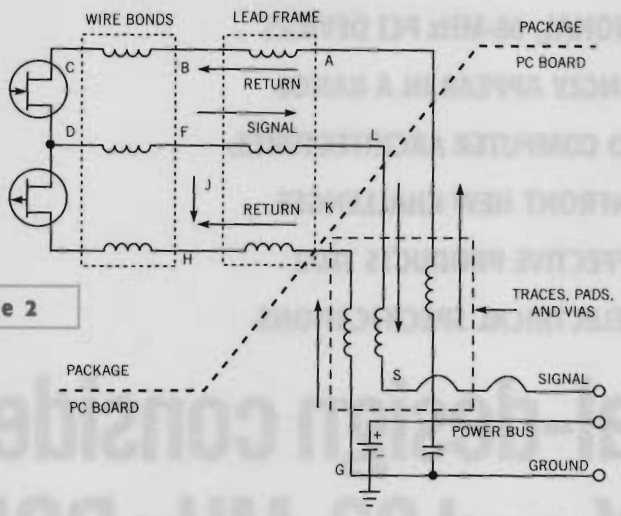

Figure 2

A lead-frame package includes self- and mutual-inductance elements. The letters correspond to equivalent points in Figure 1.

So how close do signal and return need to be? In a multilayer pc board, the signal is within 0.005 or 0.010 in. of the return plane. Ideally, the same proximity between signal and return paths should continue through the package to the silicon chip. The separation between solder balls is greater than this distance, but the distance over which that discontinuity exists is short. The shorter the length of the discontinuity, the higher the frequency must be for the discontinuity to have an adverse effect. However, if signals are close to their respective return paths but the signals are proportionally closer than that to each other, coupling between signals may occur. In other words, signals must be closer to their return paths than they are to one another.

LEAD-FRAME PACKAGES

Lead-frame packages do not incorporate return planes. A quad-flatpack lead-frame package makes signal, supply-voltage, and ground connections through thin leads (**Figure 1**). Pads on the silicon die are ball-bonded to lead fingers that run to the edge of the package. Typically, these leads extend from the package and are bent to touch pads on the pc board. The leads solder to the pc board. Some BGA packages use etched lead frames and have no return planes. The problem is that such packages without ground planes have inherently high self-inductance and mutual inductance.

You can model the semiconductor-chip and lead-frame IC package as a collection of FETs and inductors (**Figure 2**). The proximity of the inductance elements in the schematic suggests the high mutual inductance that occurs in packages that use lead frames. Leads may be on 0.003-in. centers at the center of the package and may be 0.5 in. long. Mutual inductance for such packages limits the device speeds for which you can use these packages.

A lead-frame package has far greater inductance than does a package with planes to carry return currents. However, the worst characteristic of a lead-frame package is the mutual inductance between leads in the lead frame. In a design with three signal leads on one side of a return lead, a signal is driven on the top lead (**Figure 3**). The area (A3) through which the magnetic field passes is three times the loop area of the signal lead closest to the return lead. The next lead couples two-thirds of the same area and roughly two-thirds of the magnetic field. Therefore, if the first loop has 9-nH inductance, mutual inductance is approximately 2/3×9 nH, or 6 nH. If di/dt is 60 mA/2 nsec, the mutual coupling induces dV=M di/dt=180 mV for just one adjacent signal switching, where M stands for mutual inductance. Eliminating the A3 loop is equivalent to reducing the signal-to-ground ratio from 3-to-1 to 2-to-1. Similarly, the fraction of shared area between A1 and A2 is 50%, and mutual inductance would be 4.5 nH. M di/dt would decrease to 120 mV for just one signal. Although this model is simple, it helps in visualizing the problem.

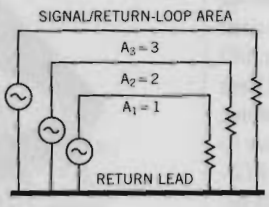

Figure 3

In a design with three signal leads on one side of a return lead, a signal is driven on the top lead. The area (A3) through which the magnetic field passes is three times that of the bottom lead.

You cannot avoid mutual inductance in packages that use a lead frame. It may be impossible to make such a package work for a given set of signals. Therefore, for high-speed ICs, you need to use packages such as BGAs that include ground and supply planes.

IDEAL RETURN-PATH DESIGN

BGA packages that have ground and supply-voltage planes for carrying signal-return currents have the potential to be extremely effective at high speeds (**Figure 4**). The current loop is for a positive-signal transition supply current GVFA, signal current FABE, displacement current j, and ground current DG. **Figure 5** is the electrically equivalent circuit for the package in **Figure 4**. The values of bond-wire inductors L_1, L_2, and L_3 would approach zero in a flip-chip package. Inductors L_4 and L_5 represent via and solder-ball connections between the package ground plane and the pc board. Inductor L_6 represents via and solder-ball inductance from the package signal trace to the pc board. Inductors L_7 and L_8 represent via and solder-ball inductance from V_{DD} in the package to the pc board.

Capacitors C_1 and C_2 represent the impedance of the power-ground plane structure.

When transistor Q_1 switches on, current is injected into Point B, and the signal begins propagating down the path toward J. As the signal propagates, a displacement current, j, immediately flows across the dielectric separating the signal from its return path, and a return current begins flowing in the return path. The displacement current, j, flows from the signal to the ground plane, which would be the case if the signal were spaced closest to the ground plane in the package.

The current loop completes through L_4, C_1, L_8, L_1, Q_1, L_2, and the displacement current. The charge between E and D is established as quickly as current can be forced through the path.

A design without return-plane connectors has the same BGA as the one in **Figure 5** but without a ground connection to the pc board at the center of the package (**Figure 6**). Without a connection from D to G, fields must propagate another way. It may seem that the width of the planes would create low impedance between planes DH and GG, but the planes are fairly far apart and there is delay associated with closing the current loop through this path. Delay translates to inductance that causes transients on V_{DD} and ground. All signals sharing this return path share these inductive effects, or ground bounce.

You can make a similar argument for signals spaced more closely to the supply plane in the BGA package when driven low by transistor Q_2. If return current is carried on the voltage plane, supply-voltage connections must be at the center of the package to avoid breaking the instantaneous-return path.

To solve this problem, you must provide both ground and power connections at the center of the package.

Figure 4

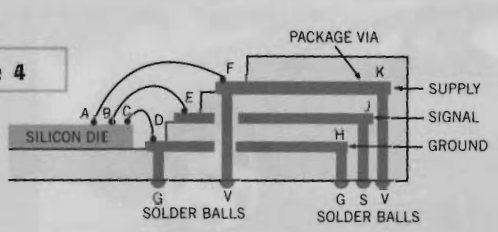

An ideal BGA package has V_{DD} and ground connections at the center of the package and among the signal-to-board connections.

Figure 5

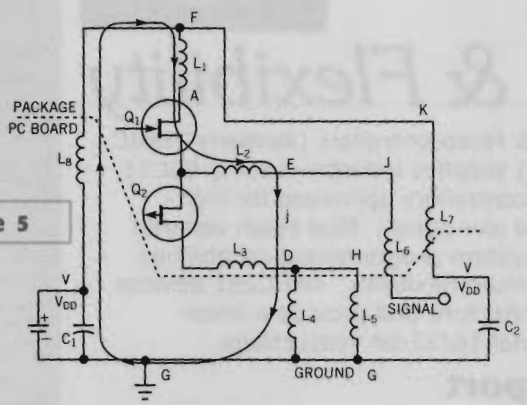

The current loop for a positive signal transition comprises supply current GVFA, signal current FABE, displacement current j, and ground current DG.

Figure 6

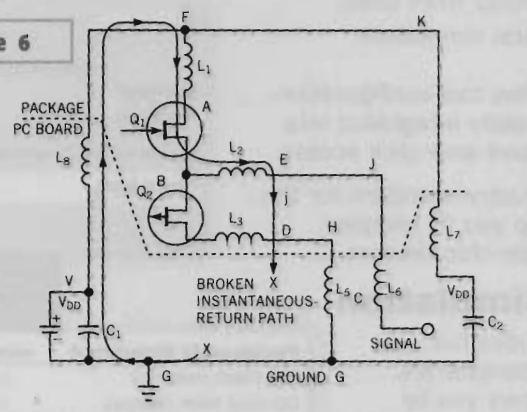

This BGA is the same as the one in Figure 5 but without a ground connection to the pc board at the center of the package and with only a V_{DD} connection at the center of the package.

Figure 7 represents a BGA package similar to the one in **Figure 5**. However, in this case, the signal path is spaced more closely to the supply-voltage plane, making the voltage plane the return path, and the voltage plane is not properly connected to the pc board among the signal pins. The signal injected at Point E is returned on Plane FK (V_{DD}). There is no connection from Point K to board V_{DD} at the signal-pin connection, S. As a result, the return current component on FK must "invent" its own return path.

The effect is the same as that of a missing return connection at the center of the package: delay and consequent inductive effects. Many signals share the path, resulting in significant ground or supply bounce or crosstalk. A similar problem exists for signals spaced more closely to the ground plane in the BGA package if you do not provide connections from the board to ground near signal connections.

To solve this problem, you must locate connections to the return-current planes among signal pins. In most cases, this approach requires both V_{DD} and ground connection. Packages in which all signals are referenced to the ground plane—that is, the signal layers are adjacent only to ground planes within the package—require only ground connections.

DESIGNING WITH NEW PACKAGING

There is considerable risk in developing and using new ASIC packages. Therefore, ASIC developers tend to build upon previous designs. Designers often don't understand existing packaging problems, so the new design rests on uncertain footing. You can use analytic tools to overcome this problem. However, these tools are difficult to use, and designers usually don't attempt to verify the results until delivery of the real packages. At that point, time-to-market pressures dominate all decisions. All too often, designers have insufficient time to fully analyze problems, so they correct only the worst problems.

You can use three approaches to reduce package-design risk. First, package design must incorporate pairing of signals with their individual return paths. Second, you should implement software modeling with field solvers and simulations. Third, you should verify software

modeling with measurements.

Even with these approaches, designers must understand the software and the electromagnetic fields, or these approaches will yield results of limited value. Lack of knowledge of these factors may lead you to overlook elements or effects in hardware, so the software model may be fundamentally incorrect.

One way to confirm software modeling is to build scale models of the package or sections of the package. When you have tested a scale model, it becomes a check on the software. In addition, working with the hardware may enable you to understand and include other effects in software analysis. Many techniques are available to simplify the task of building and effectively using a scale model. First, you do not need to build the whole package. Rather, you should divide the structure into symmetric sections. A one-quarter or one-eighth pie section (45°) of a package may include every effect you need. The planes of symmetry are boundaries across which no current flows.

Figure 7

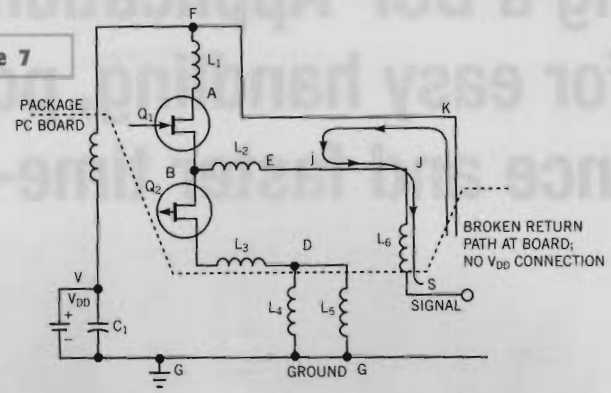

This BGA package is similar to the one in Figure 5. However, in this case, the signal path is spaced more closely to the supply-voltage plane, making the voltage plane the return path, and the voltage plane is not properly connected to the pc board among the signal pins.

Second, you should build models that are 10 or 20 times the actual size. At this scale, a patient technician can build a valuable model with a scalpel, a straight edge, pc-board material, a few transistors, and a few resistors for terminations. Time, inductance, and capacitance scale linearly, so you can replace a fast device with a transistor that is 10 or 20 times slower.

An active IC has many drive sources. Duplicating these sources would be difficult. Instead, you can use a single source with resistors to manifold the signal to the sites of the numerous drivers. The source supplies voltage adjusted to overcome the voltage drop through the resistors. The resistors equalize current to match the signals at the driver sites. The resistor values should be large enough to mask inductive effects.

You can test and correlate the physical model with software models and simulations. When the two techniques correlate, you can use the software with confidence. It is easier to try different possibilities with software, but verification by an independent method is extremely valuable. This technique has another major advantage: It directly measures package performance. Measurements on a scaled model in the time domain relate directly to performance requirements and require no inferences. In addition, it is often easier to test signal propagation in the scaled model than in the actual package. An important advantage of scaled models is that you can test them before package hardware is available and before you invest in hardware and tooling.

Figure 8

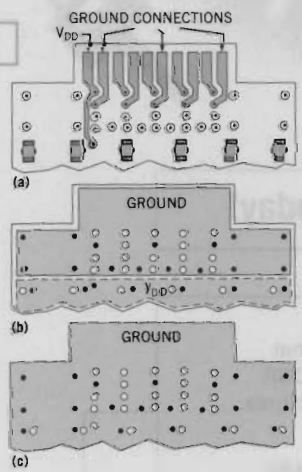

The top layer shows a set of capacitors on low-inductance footprints along the lower edge (a). In the V_{DD}, split-plane view, these capacitors connect to the V_{DD} plane segment (b). The plane segment above the split is grounded at the solid black dots. The ground-plane layer connects at the black dots to the capacitor grounds, to the ground segment, and to the connector ground pins (c).

PC-BOARD LAYOUT

You should consider four primary factors when laying out the pc boards: return-plane splits; board interfaces to connectors and IC packages; crosstalk between adjacent signals; and power-bus stabilization, or "decoupling." Meeting the requirements of the PCI-X and conventional 66-MHz PCI specifications generally requires understanding and implementing methods for dealing with these issues.

Conductors near a signal line carry return current. As a result, ground and V_{DD} planes provide return paths for the signals in the adjacent planes and return some current for the next adjacent layer. If the return current path is interrupted, ground bounce and crosstalk effects occur, resulting in glitches on quiet lines and apparent logic transitions in either high or low states. In extreme cases, reflections at this interface can cause ringing and multiple transitions on control lines.

When clearance holes in the planes around closely spaced connector pins or vias are large enough to merge, this situation can unintentionally break planes in the pin field of a connector or IC package. In other cases, the design requires a split in a plane. For example, the density

Figure 9

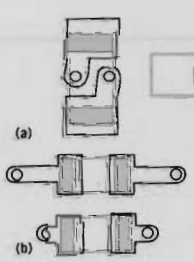

The low-inductance 0603 footprint occupies less space, uses the least expensive capacitors, and has an inductance of 0.99 nH (a). Conventional footprints for 0603 and 1206 capacitors use abnormally short traces to the vias and have inductances of 5 to 10 nH for a typical footprint and 1.8 or 2.4 for a short-trace capacitor (b).

of some designs requires that multiple supply voltages share a single plane layer with V_{DD}.

To prevent an array of clearance holes from accidentally splitting the planes, you should leave a 5-mil-wide margin between clearance holes in the planes. In boards that by design require a split plane, the best approach is to avoid having high-frequency signals cross the split. You should route high-frequency signals around the split or reference them to a plane layer that is not split.

If neither of these alternatives is possible, you must create a return path across the split. The design must incorporate an unbroken plane on another layer that continues across the split. This plane—generally a ground plane—becomes the return path across the split. The planes on both sides of the split must bypass to the continuous plane. The capacitors are most effective if the design uses low-inductance footprints. In addition, you should space the capacitors approximately 1 in. apart; occasional spacing of 1.5 inches is acceptable. If the planes on both sides of a split are capacitively coupled to the continuous plane, return current flows from the plane nearest the signal trace into the continuous plane. Return current crosses the split while in the continuous plane and returns to the plane nearest the signal trace.

CROSS CONNECTORS

If a plane carries return current to the connector, that current must pass through the connector. If it does not connect to return pins in the connector (V_{DD} or ground), the consequent interruption of return current paths causes ground bounce, glitches on quiet lines, and apparent logic transitions in either high or low states. In extreme cases, reflections at this interface can distort the primary signal, causing ringing and nonmonotonic transitions.

When assigning signals to pins of a new connector, return pins must be next to the associated signal pins. It is not enough to incorporate a large number of pins; you must also properly arrange them. However, the pinouts of many connectors, such as the PCI connector, are already assigned. The V_{DD} plane runs to the connector (**Figure 8**). Only one V_{DD} pin is at the connector, and it is poorly positioned for the signals. The return currents for signals on the layers next to the V_{DD} plane encounter an impedance discontinuity as they cross the connector. To compensate for this problem, you should split the V_{DD} plane far enough away from the connector that the edge of the V_{DD} plane can bypass to the ground plane. You should uniformly distribute connections for the ground plane among the signal pins for the full length of the connector.

LINE-TO-LINE COUPLING

If signal lines are close together and parallel over too great a length, signals from one line couple into the other. The result is spurious glitches, which cause apparent logic transitions on a quiet line in either a high or a low state. The PCI-X specification places limits on the fraction of the system noise margin that is allocated to crosstalk. These limits affect the layout of both the system board and add-in boards. Commercially available analysis software determines crosstalk between signals throughout the layout and reports instances that exceed allowable limits. When using such software, you may run lines relatively close together for short distances in the layout. Analysis shows the cumulative effect. However, if you lack access to such software, you must take a more conservative approach. The layout should minimize spacing between signals regardless of how short the distance of parallelism. If you have no analysis tools, you should space adjacent signals at least 3.3 times the distance from signal trace to the nearest return plane. Section 15 of the PCI-X specification also presents examples of six- and eight-layer pc boards with acceptable crosstalk characteristics.

The power-bus supplies voltage to the devices and must minimize voltage changes. If the voltage spikes during current steps, the logic waveforms will have supply-voltage transients. The result may be slower signal transition, reducing setup times; glitches in the waveform edges, causing timing uncertainty; or noise on the signals, reducing signal margins. In addition, voltage changes on the power bus contribute to generation of RFI.

When you space power and ground planes less than 15 mils apart, those planes bring current from capacitors all over the plane. In other words, it is unnecessary to mount capacitors on the IC pins. Bypassing some distance from the pins is effective. Compaq recommends mounting bypass capacitors on a grid pattern with spacing of roughly 1.25 in. and no greater than 1.5 in. between capacitors. Bypassing capacitors should be within 0.25 in. of corners. These requirements apply to board edges or for plane splits.

A grid arrangement of bypass capacitors has two effects: The spacing prevents RF energy from entering between planes. As a result, the system cannot launch or propagate RFI between the planes, and such a design provides better control of radiated emission from between the planes. This arrangement provides a uniform, low-impedance power source. For maximum effectiveness, minimize parasitic inductance in the decoupling capacitors, minimize trace length to vias, and do not allow capacitors to share traces or vias. Low-inductance pc-board footprints further reduce inductance of conventional surface-mount capacitors. These capacitors have vias between the solder pads and under the capacitor body, reducing the inductance associated with the body of the capacitor (**Figure 9**).

In summary, designers of PCI devices and add-in cards must strictly adhere to the PCI-X and PCI electrical specifications to develop cards that perform in all systems. These high-frequency-design methods should help achieve that goal.□

AUTHORS' BIOGRAPHIES

Joe Stoddard is a principal member of the hardware technical staff at MakeSans Systems Inc's New Business Development Division. He develops electrical-signal-integrity-analysis, power-bus-stabilization, noise-control, and pc-board hardware design. Before relocating to California, early in 2001, he was a senior member of the technical staff at Compaq (Houston), where he did similar development work for industry-standard servers.

Alan Goodrum is a staff fellow at Compaq Computer Corp, where he has worked for 15 years. In his current position, he develops technology for use in industry-standard servers, such as PCI-X and PCI Hot Plug. He is chairman of the PCI-X Workgroup and the PCI Hot-Plug Workgroup. He holds a BSEE and an MSEE from the University of Houston.